

Attorney's Docket No. P1088

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Box Patent Application
Assistant Commissioner for Patents
Washington, D.C. 20231

NEW APPLICATION TRANSMITTAL

Transmitted herewith for filing is the patent application of Inventor(s):

MATT AYERS; BEN BLACK; CHRIS BROWN; JOHN CARLSON; DAN COHN; SCOTT LAIRD;
JON MILLER; STEPHEN RAMSEY; OPHIR RONEN; PAUL SCHACHTER;
OSCAR STIFFELMAN

For (title):

METHOD AND SYSTEM FOR DIRECTING REQUESTS FOR CONTENT TO A
CONTENT SERVER BASED ON NETWORK PERFORMANCE

1. Type of Application

This new application is for a(n):


- ☒ Original (nonprovisional)
☐ Design
☐ Plant
☐ Divisional
☐ Continuation
☐ Continuation-in-part (CIP)

CERTIFICATION UNDER 37 CFR 1.10

I hereby certify that this New Application Transmittal and the documents referred to as enclosed therein are being deposited with the United States Postal Service on this date MAY 22, 2000 in an envelope as "Express Mail Post Office to Addressee" Mailing Label Number EL641402875US addressed to the: Assistant Commissioner for Patents, Washington, D.C. 20231.

John P. O'Banion

(Type or print name of person mailing paper)


(Signature of person mailing paper)

NOTE: Each paper or fee referred to as enclosed herein has the number of the "Express Mail" label placed thereon prior to mailing. 37 CFR 1.10(b).

2. Papers Enclosed Which Are Required For Filing Date Under 37 CFR 1.53(b) (Regular) or 37 CFR 1.153 (Design) Application

18 Pages of specification

15 Pages of claims

1 Pages of Abstract

5 Sheets of drawing

X formal

 informal

 The enclosed drawing(s) include photograph(s), and there is also attached a "PETITION TO ACCEPT PHOTOGRAPH(S) AS DRAWING(S)." 37 C.F.R. 1.84(b).

3. Additional papers enclosed

 Preliminary Amendment

 Information Disclosure Statement

 Form PTO - 1449

 Citations

 Authorization of Attorney(s) to Accept and Follow Instructions from Representative

 Special Comments

 Other

4. Declaration Or Oath

 Enclosed

executed by:

 inventor(s)

 legal representative of inventor(s). 37 CFR 1.42 or 1.43.

 joint inventor or person showing a proprietary interest on behalf of inventor who refused to sign or cannot be reached.

 this is the petition required by 37 CFR 1.47 and the statement required by 37 CFR 1.47 is also attached. (See item 18 below for fee.)

☐ Copy from a prior application (37 CFR 1.63(d)) (divisional or continuation only)

☒ Not Enclosed.

☒ Application is made by a person authorized under 37 CFR 1.41(c) on behalf of all of the above named inventor(s). (The declaration or oath, along with the surcharge required by 37 CFR 1.16(e) can be filed subsequently).

☐ Attached is a showing that the filing is authorized. (Not required unless called into question. 37 CFR 1.41(d)).

5. Inventorship Statement

The inventorship for all the claims in this application are:

☐ The same ☐ or

☒ Are not the same. An explanation, including the ownership of the various claims at the time the last claimed invention was made,

☐ is submitted.

☒ will be submitted.

6. Language

☒ English

☐ non-English

☐ the attached translation is a verified translation. 37 CFR 1.52(d).

7. Assignment

☒ An assignment of the invention to: INTERNAP NETWORK SERVICES

☐ is attached. A separate "ASSIGNMENT COVER LETTER ACCOMPANYING NEW PATENT APPLICATION" is also attached.

☒ will follow.

8. Benefit of Prior U.S. Application(s) (35 U.S.C. 119(e), 120 or 121)

NOTE: "In order for an application to claim the benefit of a prior filed copending national application, the prior application must name as an inventor at least one inventor named in the later filed application and disclose the named inventor's invention claimed in at least one claim of the later filed application in the manner provided by the first paragraph of 35 U.S.C. 112." 37 CFR 1.78(a).

NOTE: "In addition, the prior application must be (1) complete as set forth in § 1.51, or (2) entitled to a filing date as set forth in § 1.53(b) and include the basic filing fee set forth in § 1.16, or (3) entitled to a filing date as set forth in § 1.53(b) and have paid therein the processing and retention fee set forth in § 1.21(l) within the time set forth in § 1.53(d)." 37 CFR 1.78(a).

NOTE: "Any nonprovisional application claiming the benefit of one or more prior filed copending provisional applications must contain or be amended to contain in the first sentence of the specification following the title a reference to each such prior provisional application, identifying it as a provisional application, and including the provisional application number (consisting of the series code and serial number) and filing date." 37 CFR 1.78(a)(4).

NOTE: "Any nonprovisional application claiming the benefit of one or more prior filed copending nonprovisional applications or international applications designating the United States of America must contain or be amended to contain in the first sentence of the specification following the title a reference to each such prior application, identifying it by application number (consisting of the series code and serial number) and filing date or international application number and international filing date and indicating the relationship of the applications. Cross-references to other related applications may be made where appropriate. (See §1.14(b))." 37 CFR 1.78(2).

— Applicant(s) hereby claim(s) the benefit of the filing date of prior U.S. Application Serial No. filed on _____.

(a) Application History (title as originally filed and as last amended, serial number, and filing date of all prior applications):

Title:

Ser. No.:

Filed:

(b) Name of applicant(s) (as originally filed and as last amended), and current correspondence address of applicant(s):

Name:

Address:

NOTE: The proper reference to a prior filed PCT application which entered the U.S. national phase is the U.S. serial number and the filing date of the PCT application which designated the U.S.

NOTE: (1) Where the application being transmitted adds subject matter to the International Application then the filing can be as a continuation-in-part or (2) it is desired to do so for other reasons, then the filing can be as a continuation

NOTE: The deadline for entering the national phase in the U.S. for an international application was clarified in the Notice of April 28, 1987 (1079 O.G. 32 to 46) as follows:

"The Patent and Trademark Office considers the international application to be pending until the 22nd month from the priority date if the United States has been designated and no Demand for International Preliminary Examination has been filed prior to the expiration of the 19th month from the priority date and until the 32nd month from the priority date if a Demand for International Preliminary Examination which elected the United States of America has been filed prior to the expiration of the 19th month from the priority date, provided that a copy of the international application has been communicated to the Patent and Trademark Office within the 20 or 30 month period respectively. If a copy of the international application has not been communicated to the Patent and Trademark Office within the 20 or 30 month period, respectively, the international application becomes abandoned as to the United States 20 or 30 months from the priority date, respectively. These periods have been placed in the rules as paragraph (h) of § 1.494 and paragraph (i) of § 1.495. A continuing application under 35 U.S.C. 365(c) and 120 may be filed anytime during the pendency of the international application."

9. Priority Claim for Prior Application (35 U.S.C. 119)

_____ The prior U.S. application(s), including any prior International Application designating the U.S. identified above in item 8, in turn itself claim(s) foreign priority (ies) as follows:

(country) (appln. no.) (filed on)

(country) (appln. no.) (filed on)

(country) (appln. no.) (filed on)

The certified copy (ies)

_____ is (are) attached.

_____ has (have) been filed on _____ in prior application serial number _____ which was filed on _____.

_____ will follow.

WARNING: The certified copy of the priority application which may have been communicated to the PTO by the International Bureau may not be relied on without the need to file a certified copy of the priority application in a continuing application. This is so because the certified copy of the priority application communicated by the International Bureau is placed in a folder and is not assigned a U.S. serial number unless the national stage is entered. Such folders are disposed of if the national stage is not entered. Therefore, such certified copies may not be available if needed later in the prosecution of a continuing application. An alternative would be to physically remove the priority documents from the folders and transfer them to the continuing application. The resources required to request transfer, retrieve the folders, make suitable record notations, transfer the certified copies, enter and make a record of such copies in the continuing application are substantial. Accordingly, the priority documents in folders of international applications which have not entered the national stage may not be relied on. Notice of April 28, 1987 (1079 O.G. 32 to 46).

10. Further Inventorship Statement Where Benefit of Prior Application(s) Claimed

NOTE: "If the continuation, continuation-in-part, or divisional application is filed by less than all the inventors named in the prior application, a statement must accompany the application when filed requesting deletion of the names of the person or persons who are not inventors of the invention being claimed in the continuation, continuation-in-part, or divisional application." 37 CFR 1.62(a) [emphasis added] (dealing with the file wrapper continuation situation).

NOTE: "In the case of a continuation-in-part application which adds and claims additional disclosure by amendment, an oath or declaration as required by § 1.63 must be filed. In those situations where a new oath or declaration is required due to additional subject matter being claimed, additional inventors may be named in the continuing application. In a continuation or divisional application which discloses and claims only subject matter disclosed in a prior application, no additional oath or declaration is required and the application must name as inventors the same or less than all the inventors in the prior application." 37 CFR 1.60(c). (dealing with the continuation situation).

(complete applicable item (a) or (b) below)

- (a) ☐ This application discloses and claims only subject matter disclosed in the prior application whose particulars are set out above and the inventor(s) in this application are

☐ the same

☐ less than those named in the prior application and it is requested that the following inventor(s) identified above for the prior application be deleted:

Name:

Name:

Name:

- (b) ☐ This application discloses and claims additional disclosure and a new declaration or oath is being filed. With respect to the prior application whose particulars are set out above, the inventors in this application are

☐ the same

☐ add the following inventors

Name:

Name:

Name:

11. Maintenance of Copendency of Prior Application

NOTE: The PTO finds it useful if a copy of the petition filed in the prior application extending the term for response is filed with the papers constituting the filing of the continuation application. Notice of November 5, 1985 (1060 O.G. 27).

☐ Extension of time in prior application

(This item must be completed and the necessary papers filed in the prior application if the period set in the prior application has run)

☐ A petition, fee and response has been filed to extend the term in the prior application until _____.

☐ A copy of the petition for extension of time in the prior application is attached.

(complete this item and file conditional petition in prior application if previous item not applicable)

☐ Conditional Petition For Extension Of Time In Prior Application

☐ A conditional petition for extension of time is being filed in the pending prior application.

12. Abandonment of Prior Application (if applicable)

____ Please abandon the prior application at a time while the prior application is pending or when the petition for extension of time or to revive in that application is granted and when this application is granted a filing date so as to make this application copending with said prior application.

NOTE: According to the Notice of May 13, 1983, (103, TMOG 6-7), the filing of a continuation or continuation-in-part application is a proper response with respect to a petition for extension of time or a petition to revive and should include the express abandonment of the prior application conditioned upon the granting of the petition and the granting of a filing date to the continuing application.

NOTE: "A registered attorney or agent acting under the provisions of § 1.34(a), or of record, may also expressly abandon a prior application as of the filing date granted to a continuing application when filing such a continuing application." 37 CFR 1.138.

13. Petition For Suspension Of Prosecution For The Time Necessary To File An Amendment (if applicable)

WARNING: "The claims of a new application may be finally rejected in the first Office Action in those situations where (1) the new application is a continuing application of, or a substitute for, an earlier application, and (2) all the claims of the new application (a) are drawn to the same invention claimed in the earlier application, and (b) would have been properly rejected on the grounds of art of record in the next Office Action if they had been entered in the earlier application." MPEP § 706.07(b).

NOTE: Where it is possible that the claims on file will give rise to a first action final for this continuation application and for some reason an amendment cannot be filed promptly (e.g., experimental data is being gathered) it may be desirable to file a petition for suspension of prosecution for the time necessary.

(check the next item, if applicable)

____ There is provided herewith a Petition to Suspend Prosecution For The Time Necessary To File An Amendment (New Application Filed Concurrently)

14. Notification in Parent Application of this Filing (if applicable)

____ A notification of the filing of this application is being filed in the parent application from which this application claims priority under 35 U.S.C. 120.

15. Fee Calculation (37 CFR 1.16)

A. X Regular Application

CLAIMS AS FILED									
		Number filed		Number Extra		Rate		Basic Fee \$ 690.00	
Total									
Claims 37 CFR 1.16(c)		64	- 20	=	44	X	\$18.00	=	792.00
Independent									
Claims (37 CFR 1.16(b))		10	- 3	=	7	X	\$78.00	=	546.00
Multiple dependent claim(s), if any (37 CFR 1.16(d))						+	\$260.00	=	

- ☐ Amendment canceling extra claims enclosed.
☐ Amendment deleting multiple-dependencies enclosed.
☐ Fee for extra claims is not being paid at this time.

Filing Fee Calculation \$ 2,028.00

B. Design application
(\$310.00 - 37 CFR 1.16(f))

Filing Fee Calculation \$

C. Plant application
(\$480.00 - 37 CFR 1.16(g))

Filing Fee Calculation \$

16. Small Entity Statement(s)

X Verified Statement(s) that this is a filing by a small entity under 37 CFR 1.9 and 1.27

 is(are) attached.

X will follow.

 Status as a small entity was claimed in prior application serial number filed on , from which benefit is being claimed for this application under 35 U.S.C. 119(e), 120, 121 or 365(c) and which status as a small entity is still proper and desired. A copy of the verified statement in the prior application is included.

Filing Fee Calculation (50% of **A**, **B** or **C** above) \$ 1,014.00

17. **Request for International-Type Search (37 CFR 1.104(d))**

☐ Please prepare an international-type search report for this application at the time when national examination on the merits takes place.

18. **Fee Payment Being Made At This Time**

☒ Not Enclosed

☒ No filing fee is to be paid at this time. (This and the surcharge required by 37 CFR 1.16(e) can/will be paid subsequently.)

☐ Enclosed

☐ basic filing fee \$ _____

☐ recording assignment (\$40.00; 37 CFR 1.21(h)) \$ _____

☐ petition fee for filing by other than all the inventors or person on behalf of the inventor where inventor refused to sign or cannot be reached. (\$130.00; 37 CFR 1.47 and 1.17(h)) \$ _____

☐ for processing an application with a specification in a non-English language. (\$130.00; 37 CFR 1.52(d) and 1.17(k)) \$ _____

☐ processing and retention fee (\$130.00; 37 CFR 1.53(d) and 1.21(l)) \$ _____

☐ fee for international-type search report. (\$40.00; 37 CFR 1.21(e)) \$ _____

Total Fees Enclosed \$ _____

19. **Method of Payment of Fees**

☐ Check in the amount of \$ _____

☐ Charge Account No. _____ in the amount of \$ _____.
A duplicate of this transmittal is attached.

20. **Authorization to Charge Additional Fees**

☐ The Commissioner is hereby authorized to charge the following additional fees by this paper and during the entire pendency of this application to Account No. _____;

☐ 37 CFR 1.16(a), (f) or (g) (filing fees)

☐ 37 CFR 1.16(b), (c) and (d) (presentation of extra claims)

- ☐ 37 CFR 1.16(e) (surcharge for filing the basic filing fee and/or declaration on a date later than the filing date of the application)
- ☐ 37 CFR 1.18 (application processing fees)
- ☐ 37 CFR 1.18 (issue fee at or before mailing of Notice of Allowance, pursuant to 37 CFR 1.311(b))

21. Instructions As To Overpayment

☐ credit Account No. _____
☒ refund

22. Incorporation By Reference of Papers Identified Herein


Applicant(s) hereby incorporate(s) by reference all papers which are identified in this New Application Transmittal.

23. Correspondence Address

Please use the following correspondence address for all communications:

LaRIVIERE, GRUBMAN & PAYNE LLP
P.O. BOX 3140
MONTEREY, CA 93942
(831) 649-8800

Dated: 5/22/2000.


SIGNATURE OF ATTORNEY
John P. O'Banion, Reg. No. 33,201

CERTIFICATION UNDER 37 CFR 1.10

I hereby certify that this New Application Transmittal and the documents referred to as enclosed therein are being deposited with the United States Postal Service on this date MAY 22, 2000 in an envelope as "Express Mail Post Office to Addressee" Mailing Label Number EL641402875US addressed to the: Assistant Commissioner for Patents, Washington, D.C. 20231.

JOHN P. O'BANION

(Type or print name of person mailing paper)

(Signature of person mailing paper)

TITLE OF THE INVENTION

**METHOD AND SYSTEM FOR DIRECTING REQUESTS FOR CONTENT TO A
CONTENT SERVER BASED ON NETWORK PERFORMANCE**

CROSS-REFERENCE TO RELATED APPLICATIONS

Not Applicable

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH
OR DEVELOPMENT

Not Applicable

REFERENCE TO A MICROFICHE APPENDIX

Not Applicable

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention pertains generally to network communications, and more particularly to redirecting requests for content to a particular content server in a group of content servers based on a cost function indicative of operational characteristics of the network, such as network performance.

2. Description of the Background Art

One of the most common complaints of Internet users is that it simply takes too much time to access information. This is a particular concern for online vendors conducting e-commerce transactions. Many customers or potential customers may not return to a site if access is too slow, their browser locks up in

the middle of a transaction, transactions are duplicated because they resubmit information as a result of believing that their browser has locked up during a transaction, or they are otherwise dissatisfied with the overall speed of accessing the site. While some of the problems experienced can be caused by the site itself,
5 many of the problems experienced today are due to network performance.

One of the causes of unacceptable network performance is due to the fact that the Internet is a network of interconnected networks. As a result, more than one path may exist between the location at which the data originates and the destination to which the data will be delivered. Typically data is transmitted in
10 "packets" which include a destination address. However, there is no guarantee that the route assigned to the data to reach that destination address will be the most efficient at all times due to the dynamic nature of network capability. For example, a path that is free at one point in time may become congested at another point in time. Or, a portion of the network may become unavailable, thus causing
15 the packets to be re-routed after leaving the point of origin. Further, a network carrier who provides a portion of the intermediate path send the packets to another network carrier through a network access point if the destination is not a customer of that carrier, thereby resulting in data loss or additional re-routing. As can be seen, therefore, the packets must hop from network to network to reach
20 their destination. The result is that each hop can contribute to additional delivery delay, or even cause a loss of data. The delivery delay that may be experienced can have a detrimental effect on the data transmission, and can be analogized to the decibel (db) loss experienced when signals are transmitted through a lossy medium.

25 Rather than pursue routing solutions, a number of companies are pursuing "content delivery" techniques in an attempt to enhance performance on the Internet. These content delivery techniques include caching, intelligent routing, and content optimizing, as described below.

For example, when an internet browser sends a request for a Web page, it
30 first receives a hypertext markup language (HTML) file from the site that describes the Web page. The browser then sends one or more requests for the embedded files displayed on the Web page, such as images, audio clips, or video; that is the content. As can be seen, this process requires two or more requests for

information to be made from the user to the Web page, thus having a direct impact on delivery time. To reduce the delivery time, some "content providers" cache the site's content on multiple servers at different geographical (physical) locations. Then, when a browser sends a request for a Web page, the HTML file that is sent
5 back is accompanied by an internet protocol (IP) address that directs the browser to the geographically or topologically closest server that caches the content. The theory behind this approach is that the path to the server that is geographically or topologically closest to the user will be the fastest path.

A refinement of the foregoing approach has been to optimize content and
10 route the content based on demand. For example, the content of a Web page can be separated into "objects" representing images, video, text, and audio which are then ranked by user demand. Based on this ranking, the objects are cached at servers that are geographically closest to the users demanding access to the objects. To do so, a real-time map of Internet traffic is maintained for determining
15 to which servers to send the objects for caching based on demand.

While the foregoing techniques are designed to speed delivery of content to users, they all suffer from a common drawback; namely, that the path to the geographically closest location to the user may not be the optimum path. To the contrary, a path to an entirely different server that may not be the geographically
20 or topologically closest may provide for a faster response time. Furthermore, such techniques do not direct a user to a particular content server based on a cost function in a closed loop system. Therefore, a need exists for a method and system for providing a user with the IP address of a content server based on criteria other than geography or topology alone. The present invention satisfies
25 that need, as well as others, and overcomes the deficiencies in current techniques.

BRIEF SUMMARY OF THE INVENTION

The present invention generally comprises a method and system for pointing a user requesting access to content on a particular content server to one
30 of a plurality of content servers that may provide that content. More particularly, the content server to which the user is directed is selected based on one or more criteria that satisfy a cost function indicative of operational characteristics of the network.. The invention operates in a closed loop where additional cost

measurements are made each time a user is directed to a content server and that information becomes part of the data store for evaluating which content server should be selected for serving subsequent requests for content.

For example, the user may be directed to a particular content server if the user can access that server more quickly than another content server. While that content server may also be closest to the user in physical space, the two locations do not necessarily have to coincide. Another example would be to direct the user to a content server based on one or more criteria that are indicative of best performance for that user. Another example would be to use reliability, such as least packet loss, to the content server as a performance criteria. Another example would be to direct the user to a particular content server based on economic or financial cost criteria. Still another example would be to use the geographical location of the content server as a performance criteria.

In other words, the invention directs a user to a content server based on a cost function that takes into account one or more cost criteria indicative of operational characteristics of the network.. In this regard, note that the routing to a particular location is not being changed; routing is handled in a different manner through techniques described in, for example, U.S. Patent No. 6,009,081 which is incorporated herein by reference. In the present invention, the user is being directed to an IP address of a particular content server.

By way of example, and not of limitation, the IP address of the client who is requesting the content and the ID of the customer whose content is being requested are determined. The customer's ID is mapped to one or more content servers where the customer's content is located. Then the client's IP address is mapped to one of those content servers based on performance data to ensure that the best performance is achieved. It is important to note that, while the best performance to a content server for a particular user is one measure of cost, it is not the only measure of cost that can be used to direct the user. Nor is best performance necessarily synonymous with the geographically or topologically closest content server. Instead, the content server is selected based on one or more criteria that may include an estimate of the best performance, such as fastest response, and which may or may not translate to the geographically or topologically closest content server. Based on one or more cost factors such as

performance measurements, the client's IP address, the customer's ID, and the locations of the content servers, the present invention returns the IP address of the content server satisfying the cost function for that client. According to an optional aspect of the invention, all clients using the same DNS server would be
5 treated as satisfying the same criteria for ease of implementation.

Further objects and advantages of the invention will be brought out in the following portions of the specification, wherein the detailed description is for the purpose of fully disclosing preferred embodiments of the invention without placing limitations thereon.

10 BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be more fully understood by reference to the following drawings which are for illustrative purposes only:

FIG. 1 is a system and data flow diagram showing configuration and operation of an embodiment of the present invention.

15 FIG. 2 is a diagram depicting a geographical inference method according to the present invention.

FIG. 3 is a diagram depicting a circular intersection locating technique according to the present invention.

FIG. 4A through FIG. 4F are diagrams depicting an alternative embodiment
20 of the circular intersection technique shown in FIG. 3 which uses linear approximations.

FIG. 10 is a diagram depicting an intelligent inference technique according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

25 Referring first to FIG. 1, a schematic diagram of a network configuration 10 according to the invention is shown in the context of the Internet. FIG. 1 also depicts the preferred flow of information through the Internet according to the present invention. It will be appreciated, however, that the present invention is applicable to any interconnected network in which there are at least two servers
30 within the network that contain the same content for a site that is being requested by a client. In the example that will be described, two content servers and three private network access points (PNAP) are shown. However, the invention can be applied to any number of content servers and PNAPs. It will further be

appreciated that the invention is not limited to use in connection with PNAPs; to the contrary, the invention is applicable to any network or networks to which at least two content servers are connected. The invention is described in the context of PNAPs as part of the preferred system in which the invention can be
5 implemented.

The premise of the present invention is that a client 12 desires to access content associated with a customer's Web site 14. By "client" we mean a computer that accesses a network such as the Internet through a connection, such as a PNAP. By "customer" we mean a network server, such as a Web site
10 server, with which there is associated two or more content servers that may provide the desired content. In addition, by "content server" we mean a server with a unique internet protocol (IP) address, a group of servers with a common IP address, multiple data servers or the like, wherein the server(s) can be configured such that selection can be made from among two or more server addresses.

In the example described herein, which represents an embodiment of the system configuration in which the invention would be used, it is assumed that there is at least one PNAP, and preferably two PNAPs, connected to the Internet. It is also assumed that addressing is IP, that we are using HTTP and DNS, and that the customer will have two or more content servers connected to the same or
20 different PNAPs. For example, one content server could be connected to a first PNAP, a second content server could be connected to a second PNAP, a third content server could be connected to a third PNAP, and so on. For purposes of this example, also assume that the IP address of client 12 is 10.0.0.1, that the uniform resource locator (URL) of Web site 14 is www.superperformance.com,
25 that the IP address of the customer is 197.172.20.14, that the customer's identity (ID) is "superperformance", and that the client desires to access the home page www.superperformance.com which is "index.html". Again, however, the invention is applicable to any network configuration with multiple content servers, as well as to any IP service or addressing protocol.

In the example shown, client 12 will enter www.superperformance.com as
30 the URL in their browser. At step 100, the client's browser sends a request for www.superperformance.com who, as stated before, is assumed to be a customer with content than can be accessed on two or more content servers located within

one or more PNAPs. The request is received by a domain name system (DNS) server 16 that is maintained either by the client or by the client's internet service provider (ISP). At step 102, DNS server 16 returns 197.172.20.14 as the IP address for superperformance.com. At step 104, the client's browser sends a full
5 page request to 197.172.20.14 which connects with the customer's Web site 14.

At step 106, Web site 14 returns the full html page, which describes the embedded content as well as a URL that includes a redirector host name. The redirector host name is a new protocol that is unique to the present invention, and which contains certain information that is used for determining which content
10 server the client should access. In this example, the redirector host name would have the following basic format: client-ip.customer-id.lookup.pnap.net where client IP would correspond to 10.0.0.1, customer-id (e.g., customer identity) would correspond to "superperformance", and lookup.pnap.net would correspond to the address of a redirector DNS bridge (which is also a DNS server) which will be
15 described below. Web site 14, which will know the client's IP address, will generate the redirector host name as part of the present invention.

The more complete form of the URL for the content in this example would be http://10-0-0-1.superperformance.lookup.pnap.net/index.html. It will be appreciated that other formats conveying the desired information could be used as
20 well without departing from the present invention. For convenience, however, we have chosen to build a new layer on the existing DNS protocol. In other words, we have developed a new use of DNS. Instead of conventional IP to name mapping, we are creating an imaginary host name in a format that allows us to use DNS as the transport mechanism. The host name conveys additional information
25 that is used for accessing the redirector DNS bridge and obtaining the IP address of the content server which provides best performance to the user.

With this in mind, at step 108 the client's browser then sends the redirector host name to DNS server 16 which will not recognize the redirector host name. DNS server 16 will know, however, to make a recursive query to redirector DNS
30 bridge 18 at step 110 because "lookup.pnap.net" will be contacted as a matter of the normal DNS lookup mechanism.

At step 112 redirector DNS bridge 18 will send the client's IP address 10.0.0.1 and the customer ID "superperformance" to redirector 20. At step 114,

using information in a dynamic database 22 or an alternative data store (including DNS bridge 18 itself), redirector 20 returns the IP address for the content server based on the client's IP address, the customer's ID, and cost information associated with accessing content servers containing the requested content. In the example shown, there are two content servers 24, 26 which replicate the requested content. At step 116, redirector DNS bridge 18 sends the IP address of content server 22 to DNS server 16. At step 118, DNS server 16 sends the IP address of content server 24 to client 12. At step 120, the client's browser then contacts content server 24 with the request to serve up the embedded content, and the requested content is sent to client 12 at step 122.

Note also that FIG. 1 shows two "spanners" 28, 30 associated with content servers 24, 26 respectively. Spanners 28, 30, which in this example are part of PNAPs 34, 36, respectively, monitor and collect the cost associated with the client's request to the content server (e.g., latency, performance, etc.) The cost data is used to determine which content server will be selected by redirector 20 in response to subsequent requests for content. In the example thus described, this would take place after the client's browser contacts content server 24 and before content server 24 serves up the requested content. The performance measurements can be stored in database 22 or alternative data store or even DNS bridge 18. As stated before, performance is one measure of cost which can be used to determine the content server to which the client will be directed. Other criteria can be used as well.

As can be seen, therefore, there are essentially three components to the process. One is the determination of the IP address of the client who is requesting the content and the ID of the customer whose content is being requested. The second is mapping the customer's ID to one or more content servers where the customer's content can be accessed. The third is mapping the client's IP address to a particular content server based on one or more criteria indicative of operational characteristics of the network, such as network performance. It is important to note that network performance is not necessarily synonymous with the fastest path, and that the fastest path is not necessarily synonymous with the geographically or topographically closest content server. Instead, the content server is selected based on an evaluation of cost, which may

or may not translate to the fastest path, shortest response time, least packet loss, lowest economic cost of service, or the geographically or topographically closest content server for that client. The foregoing list is not all inclusive, but is provided by way of example of various measures of cost. Based on one or more measures
5 of cost, the client's IP address, the customer's ID, and the IP addresses of the content servers, the redirector method of the present invention returns the IP address of the content server satisfying the cost function for that client.

As an alternative to using the form of the redirector host name described above, a simplified form could be used. Referring to the previously example, the
10 simplified form would be "superperformance.lookup.pnap.net". Note that the client's IP address has been dropped. In this embodiment of the invention, redirector 20 would base its decision on the IP address of DNS server 16 instead the client's IP address. In this way, all clients who are using the same DNS server would be directed to the same content server for "superperformance.com".

15 It will be appreciated, therefore, that the cost measurements should be continuously updated. It will also be appreciated that these measurements may contain information concerning, for example, latency and reliability between a client and the content servers. In the exemplary embodiment of the invention described herein, latency measurements are obtained by employing spanners 28,
20 30 for populating and updating a data store such as database 22. A "spanner" as used herein is a software process which measures and records, for example, costs such as latency between a client and a destination in the network. Note that the preferred location of a spanner is in proximity to a content server so that the spanner can monitor the performance of the client's access to the content server.
25 If a spanner were located elsewhere, then it would not be able to determine the cost characteristics of the request and transfer of content.

Referring to FIG. 1, assume for example that client 12 is connected to a first PNAP 32, and that content servers 24, 26 are connected to second 34 and third
30 36 PNAPs, respectively. In order to determine the cost associated with accessing a content server, data is collected between client 12 and the PNAPs 34, 36. For example, using the TCP/IP protocol, this is preferably done passively by watching the client's communications with destinations in a PNAP and measuring the time between the "syn ack" and the "ack" messages. The relative measurements

between PNAPs can then be used to determine which content server can be accessed by the client with the least amount of network latency. It will be appreciated, however, that other handshakes may be employed if a protocol other than TCP/IP is used or developed. It will also be appreciated that there may be
5 other suitable methods for making such measurements, including active methods such as pinging, that can be used as well, and that the invention is not limited to the foregoing passive method which is given by way of example as the preferred method.

Note also that cost measurements between all clients and all network
10 locations (e.g., PNAPs) containing content servers may not be available. For ease of implementation, one solution would be to obtain one or more cost measurements for a client using a particular DNS server and then treat all clients using that same DNS server as if the costs were equivalent. Still another approach would be to obtain one or more cost measurements and then correlate
15 those measurements with other clients who may be similarly situated (e.g., geographically). Still another approach would be to use initial randomization associated with directing a client to a particular content server. For example, initially we may direct the client to PNAP 34 and content server 24. At that point in time, a cost measurement such as latency can be made. The next time that the
20 client requests content we may force the client to go to PNAP 36 and content server 26 where a measurement has not yet been made. In this regard, it will be appreciated that complete information can be obtained by iterating the measurements over all possibilities. It will also be appreciated that other approaches to making cost measurements may be suitable as well.

25 It should be apparent that randomization may be useful for normal system operation as well as initial operation of the system. As can be seen, we essentially need up-to-date cost measurements from each cluster or set of IP addresses (rather than each IP address) to each content server (i.e., to PNAP in which a content server is located). Therefore, even though we may have a
30 measurement for PNAP 24, we may send the client to PNAP 36 instead to get a measurement. The amount of time a client is sent to a different PNAP would be arbitrarily set, such as ten percent of the time. Further, suppose that several prefixes want to access "superperformance.com" and that one thousand

customers in a prefix want to contact the site. The randomization could be applied to clusters or sets of IP addresses, such as prefixes. This allows us to strike a balance between exploration (e.g., making cost measurements) and exploitation of information (e.g., designating a content server). Otherwise, if the client is always directed to what was previously determined to be the least cost (e.g., best performing) content server for that client, we would never obtain measurements for other content servers for that client. This "overexploitation" of measurements could lead to less than desirable performance since a different content server may satisfy the cost function for that client. On the other hand, if we are always exploring (making measurements), we will never exploit the measurements.

Accordingly, every time a client accesses a content server, the client provides performance data about that content server that can be observed by a spanner to measure and record performance, as well as provides data which can be used as an estimate of performance for the cluster or set of IP addresses containing the client's IP address. Therefore, it is possible to learn how long it takes for the client (and the group of addresses) to access the destination. When the invention directs the client to the content server that gives best performance, the system also obtains information about the client. It will be appreciated that, by sending the client to a different content server part of the time, the system adapts over time.

While the preferred approach to obtaining and maintaining the cost measurement data has been described, it may not be practical to monitor the performance or other criteria from every client's IP address or every set of addresses (e.g., such as prefixes) to a PNAP due to the large number of measurements that are required. For example, there are approximately seventy thousand prefixes in use; therefore, the number of measurements that would be required would be seventy thousand times the number of PNAPs. If prefixes are close together geographically, however, an inference can be made that physically proximate prefixes will experience similar network performance. In other words, a predictive characteristic can be added to the operation of the redirector in order to reduce the number of measurements that are required. A particularly suitable technique is to employ circular intersection to identify the geographical location of clients and then group IP addresses based on correlation of locations.

For example, referring to FIG. 2, suppose that there are multiple clients with sets of IP addresses (e.g., prefixes) P1, P2 and P3 connected to a PNAP 200. Using circular intersection techniques according to the present invention, we can determine the physical location of a set of IP addresses. In the preferred embodiment of the invention, this can be accomplished by first obtaining latency measurements between a client and PNAP 200 whose latitude and longitude is known. The minimum latency measurement in the set generally represents the most accurate measurement for purposes of circular intersection according to the invention. For example, if one-thousand latency measurements are obtained, all of the measurements except for the minimum measurement would be disregarded. The minimum latency measurement can then be converted to a distance equivalent by multiplying the measurement by the speed of light and dividing by two in the case of a round trip measurement such as is obtaining by a ping. Note, however, that this computation may overshoot the actual distance to the client in most networks due to the use of optic fiber in the network and the fact that light actually travels at a lower speed in fiber. However, only an estimate needs to be made for purposes of circular intersection according to the present invention.

In the case of FIG. 2, the minimum latency measurements from PNAP1 200 to P1 and P2 are 53 ms and 66 ms, respectively, and the minimum latency measurements from PNAP 202 to P1 and P2 are 105 ms and 100 ms, respectively. Using this information, circles can then be drawn where the PNAP location is the center of the circle and the distance equivalent of the latency is the radius of the circle. In this regard, it will be appreciated that shorter return times result in smaller (i.e., tighter) circles, thus providing more precise measurements. It will also be appreciated that any test that measures the time that it takes for a signal or data to move to a location and back can be employed without departing from the invention, and that one way of making such a measurement is by sending a TCP/IP "ping" to the client from the PNAP.

Referring also to FIG. 3, by taking similar measurements from multiple PNAPs 200, 202, 204, several circles such as circles 300, 302, 304 can be drawn and the intersection of the circles will correspond to the approximate physical location of the sets of IP addresses (e.g., a prefix), although the exact location will

not be known for a number of reasons, including the fact that light travels at a slower speed through fiber. In this example, the intersection of the circles is the approximate location of P1. An inference can then be made for sets of IP addresses that are physically close, but for which no actual measurements have been made; namely, that clients will experience similar network conditions. Therefore, in the example of FIG. 2, assume that measurements were taken for P1 and P2, but not for P3. If P3 is physically close to P1 and P2, a circle of inference 206 can be drawn for P3. This assumes, however, that set P3 was located one point in time so we know that set P3 is in the circle of inference for sets P1 and P2.

It should be apparent from the foregoing discussion that the sets of IP addresses can vary in size and, in addition, that once a circle of inference is drawn around sets of IP addresses, those sets can be used for drawing new circles of inference. For example, suppose that P1, P2 and P3 are sets of one. Once the circle of inference 206 is drawn around those sets, their combination can form a new set P4 (not shown). Now, assume that P5 (not shown) falls within the circle of inference for set P4. A new inferred set can then be created, which is the combination of P4 (P1, P2 and P3) and P5. Therefore, we can create new sets and sets of sets using this method.

Note that it is possible to implement circular intersection more efficiently by employing an approximation technique. Determining the intersection of a set of circles as the number of circles increases using conventional computational techniques can become computationally complex since the relationship between time that it takes to compute the intersections and the increased number of intersections is not linear. A linear approximation, on the other hand, is the fastest way to implement circular intersection even though it is not as accurate. This can be accomplished, for example, by summarizing the previous intersections with an intersection object. The intersection object is then intersected with a subsequent circle or another object representing that circle. Each time that an intersection is determined, the intersection is summarized with a new intersection object that is used for purposes of subsequent intersections.

For example, suppose that we have three circles 300, 302 and 304 as described before. A linear approximation of the intersection can be determined as

follows. First, the intersection 306 of two of the circles, such as circles 300, 302 would be determined as shown in FIG. 4A. Next, the area of intersection can be circumscribed by an intersection object 308, which is preferably rectangular in shape as shown in FIG. 4B. The intersection object 308 should represent the smallest bounding box that encloses the intersection. Referring also to FIG. 4C, the estimated intersection is the cross-hatched 310 area inside intersection object 308. Now, when the third circle 304 is drawn, the intersection of circle 304 and intersection object 308 can be determined. In the example shown in FIG. 4D, the intersection of circle 304 and intersection object 308 is shown as the cross-hatched area 312. However, since the intersection still involves taking curvature into account, computational efficiency can still be optimized. This can be accomplished by replacing circle 304 with a rectangular-shaped object representing circle 304, such as square 314 shown in FIG. 4E which circumscribes circle 304. The intersection is shown as the cross-hatched area 316 bounded by a new rectangular-shaped intersection object 318.

Intersection object 318 would then be used for the next intersection with a circle by again replacing the circle with a rectangular-shaped object representing that circle. As can be seen, therefore, each intersection creates a new intersection object that can be used for a subsequent intersection.

In order to reduce computational complexity, note in FIG. 4E that square 314 has been aligned with intersection object 308 such that corresponding sides are parallel. This ensures that intersection object 318 is rectangular-shaped. To further reduce computational complexity, it is also possible to rotate the coordinate system of the intersecting circles so that the longitudinal axis through intersection object 308 is oriented vertically and aligned with the vertical axis in a normal Cartesian coordinate system as shown in FIG. 4F. Additional refinements include, but are not limited to:

- (1) if one circle contains another circle, discarding the larger circle to provide a tighter intersection;
- (2) if more than two circles in a set are to be intersected and the circles are similar in size, starting with the two circles that are the farthest apart;
- (3) if more than two circles in a set are to be intersected and the circles are different in size, starting with the two circles having the smallest radii to

provide the tightest intersection.

Therefore, the above-described method summarizes previous intersections with an intersection object. The intersection object is then intersected with a subsequent circle or an object representing the subsequent circle. While the intersection objects are preferably rectangular-shaped as described above, they could alternatively have other shapes that are also polygons. The method reduces computational complexity and increases computational speed by employing a linear computation technique, at the expense of accuracy, rather than employing a technique that is not linear. In other words, if n represents the number of circles being intersected, the computational complexity for exact intersection is $O(n^3)$ whereas it is $O(n)$ in the approximation method described above.

Accordingly, it can be seen that a spanner measures performance characteristics while circular intersection estimates the physical location of a client. Accordingly, circular intersection is an operational tool that can be used to speed up populating the store of cost measurements when actual measurements are not available; it is not, however, an essential part of the invention.

Referring now to FIG. 10, an "intelligent inference" method according to the invention will be described which is more accurate than drawing a circle of inference around physically proximate sets of IP address as was described previously in connection with FIG. 2. In FIG. 4, let C_a represent a first content server, let C_b represent a second content server, let S_1 , S_2 , S_3 and S_4 represent sets of IP addresses, let L represent latency, and let D represent distance. We can use a spanner to obtain latency measurements from content server C_a to S_1 (L_{1a}), S_2 (L_{2a}) and S_3 (L_{3a}), and use circular intersection to determine the distances between S_1 and S_4 (D_{14}), S_2 and S_4 (D_{24}), and S_3 and S_4 (D_{34}). With this information, we can obtain a weighted average of latency measurements to infer the latency measurement L_{4a} from content server C_a to S_4 as follows:

$$L_{4a} \approx \left(\frac{1}{D_{14}} (L_{1a}) + \frac{1}{D_{24}} (L_{2a}) + \frac{1}{D_{34}} (L_{3a}) \right) \left(\frac{1}{\frac{1}{D_{14}} + \frac{1}{D_{24}} + \frac{1}{D_{34}}} \right)$$

By taking the weighted average in this manner, we obtain an approximation of the latency from content server C_a to S_4 . Note also that inferences can be made in

other ways, including, but not limited to, the following:

1. Monitoring sets of IP addresses and then treating all IP addresses in the set in the same fashion.

2. Aggregating prefix groups based on similarities and treating all such aggregated prefix groups in the same fashion.

3. Disaggregating prefix groups into subgroups that have similar characteristics and treating the subgroups in the same fashion.

4. Placing IP addresses into sets based on some criteria of similarity such as, but not limited to, performance, geographical location, and other correlation factors, and treating all IP addresses in the sets in the same fashion.

In other words, inferences are ideally made based on sets of IP addresses. A set can be a single IP address (e.g., a set of one), a group of IP address in a BGP4 prefix if the prefix is small, the union of two sets wherein the two sets are geographically close, the union to sets wherein the two sets experience similar performance in relation to content servers, and the like.

Therefore, it will be appreciated that efficient implementation of the present invention requires cost measurements about the network that can be used to determine if directing a client to a particular content server will satisfy a cost function, and it will further be appreciated that the decision may be irrespective of the actual physical location of the client and the content server. Initially, however, cost measurements may not be available for all clients and content servers. As a result, some amount of randomization or estimation may be necessary until the system is trained and the store of performance data is populated. In addition, knowledge of the physical location of the client and/or sets of IP addresses may be helpful for drawing an inference that proximate physical locations may experience similar network conditions. Such inferences can be used as an alternative to obtaining measurements from each client IP address to each PNAP or content server. Note, however, that some measurements, such as latency, need to be updated frequently due to changes in network conditions whereas geographical locations of clients are static. Therefore, inferences should be drawn close in time to when measurements of a dynamic nature are made.

It will also be appreciated that the system of the present invention will automatically balance network load among PNAP's. Another load balancing

option would be to provide the customer with a software interface that allows them to report load on the content servers so that redirection can also be based on balancing server load.

It will be appreciated that the method and system would preferably be implemented using a programmed data processor such as a computer and associated software to carry out the functions described herein. For example, DNS bridge 18, redirector 20, and spanners 28, 30 would be implemented as software running on a programmed data processor such as a conventional computer. Those skilled in the art will appreciate that the actual code to implement the invention as described herein can be written for various platforms using conventional programming techniques. Accordingly, the details of the code are not presented herein.

As can be seen, therefore, the invention provides for obtaining cost measurements of various types associated with serving a request for content to a user and directing the user to a particular content server based in satisfaction of a cost function indicative of operational characteristics of the network. The invention operates in a closed loop wherein an additional cost measurement is made each time a user is directed to a content server and that information becomes part of the data store for evaluating which content server should be selected for serving subsequent requests for content from that content server.

Although the description above contains many specificities, these should not be construed as limiting the scope of the invention but as merely providing illustrations of some of the presently preferred embodiments of this invention. Thus the scope of this invention should be determined by the appended claims and their legal equivalents. Therefore, it will be appreciated that the scope of the present invention fully encompasses other embodiments which may become obvious to those skilled in the art, and that the scope of the present invention is accordingly to be limited by nothing other than the appended claims, in which reference to an element in the singular is not intended to mean "one and only one" unless explicitly so stated, but rather "one or more." All structural, chemical, and functional equivalents to the elements of the above-described preferred embodiment that are known to those of ordinary skill in the art are expressly incorporated herein by reference and are intended to be encompassed by the

present claims. Moreover, it is not necessary for a device or method to address each and every problem sought to be solved by the present invention, for it to be encompassed by the present claims. Furthermore, no element, component, or method step in the present disclosure is intended to be dedicated to the public
5 regardless of whether the element, component, or method step is explicitly recited in the claims. No claim element herein is to be construed under the provisions of 35 U.S.C. 112, sixth paragraph, unless the element is expressly recited using the phrase "means for."

CLAIMS

What is claimed is:

1. A method for directing a network client requesting access to content
5 to one of a plurality of content servers that can provide said content, comprising:
directing a network client to a said one of said content servers based on
one or more cost measurements indicative of operational characteristics of the
network.
- 10 2. A method as recited in claim 1, further comprising:
obtaining a new cost measurement when said network client accesses said
content server; and
using said new cost measurement as an indicator of operational
characteristics of the network in connection with subsequent requests for access
15 to said content that can be provided by said content server.
3. A method as recited in claim 1, wherein said content servers are
associated with a network server having an identity, and wherein said network
client requests content from said network server, and further comprising:
20 mapping the identity of the network server to said content servers.
4. A method as recited in claim 1, further comprising measuring
network performance between said network client and a said one of said content
servers.
- 25 5. A method as recited in claim 1, wherein an attribute of network
performance comprises network latency.
6. A method as recited in claim 5, wherein network latency is measured
30 passively by determining the time between a syn ack message sent by said
network client and an ack message sent by one of said content servers.

7. A method as recited in claim 4, further comprising measuring network performance between said network client and another of said content servers.

5 8. A method as recited in claim 1, further comprising determining the location of said network client by circular intersection and inferring network performance associated with accessing said content server to said network client based on a performance measurement to said content server in connection with another physically proximate network client.

10 9. A method as recited in claim 8, wherein said circular intersection comprises:

(a) measuring the time that it takes for data to move from a plurality of network server locations to said client;

15 (b) converting said times to distance equivalents;

(c) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and

20 (d) determining the physical location of said network client from the intersection of said circles.

10. A method as recited in claim 1, further comprising inferring network performance of serving said network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said network client and performance of said content server for said other network clients.

11. A method as recited in claim 1, further comprising:

30 (a) measuring network latency between a content server and a plurality of other network clients;

(b) determining physical distances between said other network clients and said network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said network client; and

5 (d) inferring operational characteristics associated with a plurality of network clients to said network client using said weighted average.

12. A method for directing a network client requesting access to content from a network server to one of a plurality of content servers that can provide said
10 content, each said content server having an address, said network server having an identity, said method comprising:

(a) identifying a network server associated with content requested by said network client;

15 (b) identifying a said one of said content servers based on said identity of said network server and one or more cost measurements indicative of operational characteristics of the network; and

(c) providing the network client with the address of said content server identified in step (b).

20 13. A method as recited in claim 12, further comprising:

(d) obtaining a new cost measurement when said network client accesses said content server; and

25 (e) using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

14. A method as recited in claim 12, further comprising measuring network performance between said network client and a said one of said content servers.

30

15. A method as recited in claim 12, wherein an attribute of network performance comprises network latency.

16. A method as recited in claim 15, wherein network latency is measured passively by determining the time between a syn ack message sent by said network client and an ack message sent by said one of said content servers.

5 17. A method as recited in claim 14, further comprising measuring network performance between said network client and another of said content servers.

10 18. A method as recited in claim 12, determining the location of said network client by circular intersection and inferring network performance associated with accessing said content server to said network client based on a performance measurement to said content server in connection with another physically proximate network client.

15 19. A method as recited in claim 18, wherein said circular intersection comprises:

- (a) measuring the time that it takes for data to move from a plurality of network server locations to said client;
- (b) converting said times to distance equivalents;
- 20 (c) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and
- (d) determining the physical location of said network client from the intersection of said circles.

25 20. A method as recited in claim 12, further comprising inferring network performance of serving said network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said network client and performance of said content server for said other
30 network clients.

23. A method as recited in claim 12, further comprising:

(a) measuring network latency between a content server and a plurality of other network clients;

(b) determining physical distances between said other network clients and said network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said network client; and

(d) inferring operational characteristics associated with a plurality of network clients to said network client using said weighted average.

22. A method for directing a network client requesting access to content from a network server to one of a plurality of content servers providing said content for said network server, each said content server having an address, said network server having an identity, said method comprising returning the address of a content server that may provide said content the identity of the network server based on one or more cost measurements indicative of operational characteristics of the network, obtaining a new cost measurement when said network client accesses said content server, and using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

23. A method as recited in claim 22, further comprising measuring network performance between said network client and a said one of said content server.

24. A method as recited in claim 22, wherein an attribute of network performance comprises network latency.

25. A method as recited in claim 24, wherein network latency is measured passively by determining the time between a syn ack message sent by said network client and an ack message sent by said one of said content servers.

5 26. A method as recited in claim 23, further comprising measuring network performance between said network client and another of said content servers.

10 27. A method as recited in claim 22, further comprising determining the location of said network client by circular intersection and inferring network performance associated with accessing said content server to said network client based on a performance measurement to said content server in connection with another physically proximate network client.

15 28. A method as recited in claim 27, wherein said circular intersection comprises:

- (a) measuring the time that it takes for data to move from a plurality of network server locations to said client;
- (b) converting said times to distance equivalents;
- 20 (c) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and
- (d) determining the physical location of said network client from the intersection of said circles.

25 29. A method as recited in claim 22, further comprising inferring network performance of serving said network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said network client and performance of said content server for said other
30 network clients.

30. A method as recited in claim 22, further comprising:

(a) measuring network latency between a content server and a plurality of other network clients;

5 (b) determining physical distances between said other network clients and said network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said network client; and

10 (d) inferring operational characteristics associated with a plurality of network clients to said network client using said weighted average.

31. A method for determining the physical location of a network client comprising:

15 (a) measuring the time that it takes for data to move from a plurality of network server locations to a network client;

(b) converting said times to distance equivalents;

(c) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and

20 (d) determining the physical location of said network client from the intersection of said circles.

32. A method for inferring operational characteristics associated with a plurality of network clients to an inferable network client, comprising:

25 (a) measuring network latency between a network server and a plurality of network clients;

(b) determining physical distances between said network clients and an inferable network client; and

30 (c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said network server and said inferable network client.

33. A system for directing a network client requesting access to content to one of a plurality of content servers that can provide said content, comprising:
a programmed data processor; and
5 programming associated with said programmed data processor for directing a network client to a said one of said content servers based on one or more cost measurements indicative of operational characteristics of the network.

10 34. A system as recited in claim 33, further comprising programming associated with said programmed data processor for:
obtaining a new cost measurement when said network client accesses said content server; and
using said new cost measurement as an indicator of operational
15 characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

35. A system as recited in claim 33, wherein said content servers are associated with a network server having an identity, and wherein said network
20 client requests content from said network server, and further comprising:
programming associated with said programmed data processor mapping the identity of the network server to said content servers.

36. A system as recited in claim 33, further comprising programming
25 associated with said programmed data processor for measuring network performance between said network client and a said one of said content servers.

37. A system as recited in claim 33, wherein an attribute of network
30 performance comprises network latency.

38. A system as recited in claim 37, wherein network latency is measured passively by determining the time between a syn ack message sent by said network client and an ack message sent by one of said content servers.

39. A system as recited in claim 36, further comprising programming associated with said programmed data processor for measuring network performance between said network client and another of said content servers.

5 40. A system as recited in claim 33, further comprising programming associated with said programmed data processor for determining the location of said network client by circular intersection and inferring network performance associated with accessing said content server to said network client based on a performance measurement to said content server in connection with another
10 physically proximate network client.

41. A system as recited in claim 40, wherein said circular intersection comprises:

(a) measuring the time that it takes for data to move from a plurality of
15 network server locations to said client;

(b) converting said times to distance equivalents;

(c) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and

20 (d) determining the physical location of said network client from the intersection of said circles.

42. A system as recited in claim 33, further comprising programming associated with said programmed data processor for inferring network
25 performance of serving said network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said network client and performance of said content server for said other network clients.

30

43. A system as recited in claim 33, further comprising programming associated with said programmed data processor for

(a) measuring network latency between a content server and a plurality of other network clients;

5 (b) determining physical distances between said other network clients and said network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said network client; and

10 (d) inferring operational characteristics associated with a plurality of network clients to said network client using said weighted average.

44. A system for directing a network client requesting access to content from a network server to one of a plurality of content servers that can provide said content, each said content server having an address, said network server having an identity, said method comprising:

15 (a) a programmed data processor; and

(b) programming associated with said programmed data processor for

(i) identifying a network server associated with content requested by said network client;

20 (ii) identifying a said one of said content servers based on said identity of said network server and one or more cost measurements indicative of operational characteristics of the network; and

(iii) providing the network client with the address of said content server identified in step (ii).

45. A system as recited in claim 44, further comprising programming associated with said programmed data processor for:

30 obtaining a new cost measurement when said network client accesses said content server; and

using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

46. A system as recited in claim 44, further comprising programming associated with said programmed data processor for measuring network performance between said network client and a said one of said content servers.

5

47. A system as recited in claim 46, wherein an attribute of network performance comprises network latency.

48. A system as recited in claim 47, wherein network latency is measured passively by determining the time between a syn ack message sent by said network client and an ack message sent by said one of said content servers.

49. A system as recited in claim 46, further comprising programming associated with said programmed data processor for measuring network performance between said network client and another of said content servers.

50. A system as recited in claim 44, further comprising programming associated with said programmed data processor for determining the location of said network client by circular intersection and inferring network performance associated with accessing said content server to said network client based on a performance measurement to said content server in connection with another physically proximate network client.

51. A system as recited in claim 50, wherein said circular intersection comprises:

- (a) measuring the time that it takes for data to move from a plurality of network server locations to said client;
- (b) converting said times to distance equivalents;
- (c) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and
- (d) determining the physical location of said network client from the intersection of said circles.

52. A system as recited in claim 44, further comprising programming associated with said programmed data processor for inferring network performance of serving said network client from said content server by
5 determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said network client and performance of said content server for said other network clients.

10 53. A system as recited in claim 44, further comprising programming associated with said programmed data processor for:

(a) measuring network latency between a content server and a plurality of other network clients;

15 (b) determining physical distances between said other network clients and said network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighted average comprises an estimate of the latency between said content server and said network client; and

20 (d) inferring operational characteristics associated with a plurality of network clients to said network client using said weighted average.

54. A system for directing a network client requesting access to content from a network server to one of a plurality of content servers providing said content for said network server, each said content server having an address, said
25 network server having an identity, said system comprising a programmed data processor and programming associated with said programmed data processor for returning the address of a content server that may provide said content the identity of the network server one or more cost measurements indicative of operational characteristics of the network, obtaining a new cost measurement when said
30 network client accesses said content server, and using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

55. A system as recited in claim 54, further comprising programming associated with said programmed data processor for measuring network performance between said network client and a said one of said content servers.

5

56. A system as recited in claim 55, wherein an attribute of network performance comprises network latency.

57. A system as recited in claim 56, wherein network latency is measured passively by determining the time between a syn ack message sent by said network client and an ack message sent by said one of said content servers.

10

58. A system as recited in claim 55, further comprising programming associated with said programmed data processor for measuring network performance between said network client and another of said content servers.

15

59. A system as recited in claim 54, further comprising programming associated with said programmed data processor for determining the location of said network client by circular intersection and inferring network performance associated with accessing said content server to said network client based on a performance measurement to said content server in connection with another physically proximate network client.

20

60. A method as recited in claim 59, wherein said circular intersection comprises:

25

(a) measuring the time that it takes for data to move from a plurality of network server locations to said client;

(b) converting said times to distance equivalents;

(c) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and

30

(d) determining the physical location of said network client from the intersection of said circles.

61. A system as recited in claim 54, further comprising programming associated with said programmed data processor for inferring network performance of serving said network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said network client and performance of said content server for said other network clients.

62. A system as recited in claim 54, further comprising programming associated with said programmed data processor for:

(a) measuring network latency between a content server and a plurality of other network clients;

(b) determining physical distances between said other network clients and said network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighted average comprises an estimate of the latency between said content server and said network client; and

(d) inferring operational characteristics associated with a plurality of network clients to said network client using said weighted average.

63. A system for determining the physical location of a network client comprising:

(a) a programmed data processor; and

(b) programming associated with said programmed data processor for

(i) measuring the time that it takes for data to move from a plurality of network server locations to a network client;

(ii) converting said times to distance equivalents;

(iii) forming a plurality of intersecting circles using said distance equivalents as the radius of circles with said network server locations as the center; and

(iv) determining the physical location of said network client from the intersection of said circles.

64. A system for inferring operational characteristics associated with a plurality of network clients to an inferable network client, comprising:

(a) a programmed data processor; and

(b) programming associated with said programmed data processor for

5 (i) measuring network latency between a network server and a plurality of network clients;

(ii) determining physical distances between said network clients and an inferable network client; and

10 (iii) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said network server and said inferable network client.

15

ABSTRACT OF THE DISCLOSURE

- A method and system for pointing a Internet user requesting access to content on a particular web site to one of a plurality of content servers associated with that web site based on one or more cost measurements associated with
- 5 serving said content. The identification of the web site is mapped to the content servers associated with that web site. The user's IP address is then mapped to a particular content server based on a cost function and the IP address of the selected content server is provided to the user.






FIG. 2




FIG. 3




FIG. 4A




FIG. 4B




FIG. 4C




FIG. 4D




FIG. 4E




FIG. 4F




FIG. 5